

EXAMINER'S AMENDMENT

1. An examiner's amendment to the record appears below. Should the changes and/or additions be unacceptable to applicant, an amendment may be filed as provided by 37 CFR 1.312. To ensure consideration of such an amendment, it MUST be submitted no later than the payment of the issue fee.

Authorization for this examiner's amendment was given in a telephone interview with Applicant's representative "Stephen D. LeBarron" on 04/17/2009.

Following subject matter was discussed and agreed upon during the telephonic interview conducted on 04/17/2009.

- Claims 13-14 and 33-34 have been canceled without prejudice
- All independent claims 1, 15 and 35 has amended to overcome possible 112 second rejected and language has been amended to recite, "before a storage operating system executing on the cluster partner is fully active" in [claim 1, lines 4-5], [claim 15, lines 3-4 and 11-12], [claim 35, lines 4-6 and 14-15] respectively or the term "early/early in the booting process" has been taken out completely as follows:
 - claim 1, line 4
 - claim 15, lines 3, 5 and 11
 - claim 20, lines 8
 - claim 24, line 6
 - claim 28, line 7
 - Claim 35, line 4, 7 and 14.

2. Final set of amended claims will be written as follows:

Claim 1. A method for initiating a peer-to-peer communication session, the method comprising:

Initiating a boot process;

Initializing a cluster connection manager in the booting process before a storage operating system executing on the cluster partner is fully active;

Initiating, by the cluster connection manager, a first remote direct memory access (RDMA) read operation directed to a cluster partner before a storage operating system executing on the cluster partner is fully active, the RDMA read operation bypassing the operating system;

Performing in response to a successful first RDMA read operation, a first RDMA write operation to the cluster partner;

Performing, in response to a successful RDMA write operation, a second RDMA read operation directed to the cluster partner; and

Performing, in response to a successful second RDMA read operation, a second RDMA write operation to the cluster partner before a storage operating system executing on the cluster partner is fully active, in the booting process.

Claim 2. The method of claim 1 wherein the step of attempting a first RDMA read operation further comprises the step of issuing a RDMA read operation to the cluster partner requesting a pre-set memory address location that is associated with a status variable on the cluster partner.

Claim 3. The method of claim 1 further comprising:

- exchanging a set of peer connection information;
- passing a set of client information to the cluster partner;
- creating a set of appropriate communication ports;
- alerting the cluster partner of a ready status; and
- alerting a set of clients that the cluster partner is in a ready state.

Claim 4. The method of claim 3 wherein the set of peer connection information comprises a version number.

Claim 5. The method of claim 1 wherein the step of passing a set of client information to the cluster partner further comprises:
collecting, from a set of clients, the set of client information; and
transferring the collected set of client information to the cluster partner.

Claim 6. The method of claim 5 wherein the client information comprises a number of communication ports required.

Claim 7. The method of claim 5 wherein the set of client information further comprises an amount of memory requested by a particular client.

Claim 8. The method of claim 1 wherein the cluster partner is a storage system.

Claim 9. The method of claim 1 wherein the cluster partner is an application server

Claim 15. A method comprising:

Initiating a boot process;

Initializing a cluster connection manager in the boot process before a storage operating system executing on the cluster partner is fully active;

Initiating, a peer-to-peer communication session, by a cluster connection manager, before a storage operating system executing on the cluster partner is fully active which bypasses an operating system on a storage system by attempting a first remote direct memory access read operation directed to a predefined hardware address and a predefined port number, the predefined hardware address and the predefined port number previously known to support a RDMA operation; and

Performing, in the booting process before a storage operation system executing on the cluster partner is fully active, in response to a successful initiating, a first remote direct memory access write operation directed to the predefined hardware address and the predefined port number.

Claim 16. The method of claim 15 further comprising:

Performing, in response to a successful first remote direct memory access write, a second remote direct memory access read operation directed to the predefined hardware address and the predefined port number.

Claim 17. The method of claim 15 wherein the predefined hardware address comprises a fibre channel identifier.

Claim 18. The method of claim 15 wherein the predefined port number comprises a virtual interface.

Claim 19. The method of claim 15 wherein the first remote direct memory access is delivered to a predefined memory address storing booting status information.

Claim 20. A system configured to establish reliable peer-to-peer communication among storage system of a clustered environment, the system comprising:

- a booting process executed by a processor;
- a peer process executing on each storage system partner having an operating system; and
- a cluster connection manager executing on each storage system partner, the cluster connection manager establishing a reliable peer-to-peer connection between each peer process in the booting process before a storage operating system executing on a cluster partner is fully active be connecting to a predetermined port number using a predetermined network address, the reliable peer-to-peer connection bypassing the operating system and initiate a remote direct memory access (RDMA) read operation directed to a cluster partner.

Claim 21. The system of claim 20 wherein the reliable peer-to-peer connection is established without requiring a storage operating system executing on each storage system partner to be fully functioning.

Claim 22. The system of claim 20 wherein the peer-to-peer connection is a virtual interface connection.

Claim 23. The system of claim 20 wherein the peer process is a cluster connection client that requests services from the cluster connection manager

Claim 24. A system configured to open initial peer-to-peer connection over a cluster interconnect, the system comprising:

- a storage system having an operating system;
- a booting process executed by a processor;
- a cluster connection manager executing on the storage system, the cluster connection manager configured to establish a peer connection in the booting process before a storage operating system executing on a cluster partner is fully active on a predetermined port number and using a predetermined network address within the storage system the peer-to-peer connection bypassing the operating system and initiate a remote direct memory access (RDMA) read operation directed to a cluster partner; and
- a process executing on the storage system, the process configured to use the established peer connection for communication.

Claim 25. The system of claim 24 wherein the peer-to-peer connection is a virtual interface connection.

Claim 26. The system of claim 24 wherein the process executing on the storage system is a cluster connection client that requests services from the cluster connection manager.

Claim 27. The system of claim 24 wherein the process executing on the storage system communicates with a cluster partner using the established peer connection

Claim 28. A system configured to accept the initiation of a peer-to-peer connection over a cluster interconnect, the system comprising:

 a storage system having an operating system;

 a booting process executed by a process;

 a cluster connection manager executing on the storage system, the cluster connection manager executing on the storage system, the cluster connection manager configured to accept a peer connection on a predetermined port number and using a predetermined network address within the storage system and initiate a remote direct memory access (RDMA) read operation to a cluster partner in the booting process before a storage operating system executing fully active; and

 a process executing on the storage system, the process configured to read information from the established peer connection.

Claim 29. The system of claim 28 wherein the peer-to-peer connection is a virtual interface connection.

Claim 30. The system of claim 28 wherein the process executing on the storage system is a cluster connection client that requests services from the cluster connection manager.

Claim 31. The system of claim 28 wherein the process executing on the storage system reads information from a cluster partner.

Claim 32. The system of claim 28 wherein the information comprises heartbeat signals

Claim 35. A computer readable medium containing executable program instructions executed by a processor, comprising:

Program instructions that initiate a booting process;

Program instructions that initialize a cluster connection manager in the booting process before a storage operating system executing the cluster partner is fully active;

Program instructions that initiate in the booting process, a first remote direct memory access (RDMA) read operation before a storage operating system executing on a cluster partner is fully active;

Program instructions that perform, in response to a successful RDMA write operation, a second RDMA read operation directed to the cluster partner; and

Program instructions that perform in the booting process before a storage cooperating system executing on the cluster partner is fully active, in response to a successful second RDMA read operation, a second RDMA write operation to the cluster partner.

Allowable Subject Matter

3. The following is an Examiner's statement of reasons for allowance:
 - Claims 1, 15, 20, 24, 28 and 35 are allowed.
4. Prior art singly or in combination fails to teach claim limitation, among other things, the combination of, initializing a cluster connection manager in the booting process before a storage operating system executing on the cluster partner is fully active and directing first remote RDMA read operation and in response to successful read operation follows a RDMA write operation, Further in response to first successful "RDMA" read and write operation a second set of "RDMA" read and write operation is performed.
 - Claims 2-9 are dependent from independent claim 1 above and further limit the claim.
 - Claims 16-19 depend from independent claim 15 above and further limit the claim.
 - Claims 21-23 are dependent from independent claim 20 above and further limit the claim.
 - Claims 25-27 are dependent from independent claim 24 above and further limit the claim.

- Claims 29-32 are dependent from independent claim 28 above and further limit the claim.

Any comments considered necessary by applicant must be submitted no later than the payment of the issue fee and, to avoid processing delays, should preferably accompany the issue fee. Such submission should be clearly labeled "Comments of Reasons for Allowance".

Conclusion

Any inquiry concerning this communication or earlier communications from the examiner should be directed to TAUQIR HUSSAIN whose telephone number is (571)270-1247. The examiner can normally be reached on 7:30 AM to 5:00 PM.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, John Follansbee can be reached on 571 272 3964. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

Application/Control Number: 10/622,565
Art Unit: 2452

Page 12

/T. H./
Examiner, Art Unit 2452

/Kenny S Lin/
Primary Examiner, Art Unit 2452